

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	2771498	compar\$3 nera sort\$3 near order\$1	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 11:33
L2	34	compar\$3 near sort\$3 near order\$1	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 11:33
L3	31444492	compar\$3.near sort\$3.near order\$1.and@ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 11:33
L4	23	compar\$3 near sort\$3 near order\$1 and @ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 11:40
L5	0	4 and refreshing	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 11:34
L6	2	refreshing near window and search near results	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 11:34
L7	46	refreshing same search same results	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 11:40
L8	23	7 and @ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 11:46

EAST Search History

L9	89818	((compar\$3 or examin\$3) near (result\$1 or (search near result\$1) or (sort near order\$1 or order))) and @ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 11:50
L10	89818	((compar\$3 or examin\$3) near ((result\$1 or (search near result\$1) or (sort near order\$1) or order)) and @ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 11:51
L11	89818	((compar\$3 or examin\$3) near (result\$1 or (search near result\$1) or (sort near order\$1) or order)) and @ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 11:51
L12	76738	((compar\$3) near (result\$1 or (search near result\$1) or (sort near order\$1) or order)) and @ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 13:15
L13	12906	((compar\$3) near ((search near result\$1) or (sort near order\$1) or order)) and @ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 11:52
L14	215	((compar\$3) near ((search near result\$1) or (sort near order\$1))) and @ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 11:52
L15	14	((compar\$3) near ((search near result\$1) or (sort near order\$1))) and refresh\$3 and @ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 11:59
L16	13252	((compar\$3) near ((search near result\$1) or (sort near order\$1 or order\$3))) and @ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 11:59

EAST Search History

L17	215	((compar\$3) near ((search near result\$1) or (sort near order\$1)) and @ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 12:15
L18	201	17 not 15	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 11:59
L19	83	18 and display\$3	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 11:59
L20	83	18 and (display\$3 or gui)	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 12:15
L21	0	compar\$3 adj3 sort adj3 oreer\$3	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 12:15
L22	29	compar\$3 adj3 sort adj3 order\$3	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 12:15
L23	15	22 and @ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 13:15
L24	5	sort near order same location same re-sort\$3	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 13:15

EAST Search History

L25	4	24 and @ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 13:28
L26	0	re-sort\$3 near (serach near results) and @ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 13:23
L27	4	re-sort\$3 near (search near results) and @ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 13:26
L28	0	(compar\$3 near sort near order) and re-sort\$3 near (search near results) and @ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 13:26
L29	0	(compar\$3 near sort near order) and re-sort\$3 and (search near results) and @ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 13:26
L30	0	(compar\$3 near sort near order) and re-sort\$3 and ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 13:26
L31	1054	(compar\$3 and re-sort\$3) and ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 13:27
L32	7	(compar\$3 near re-sort\$3) and ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 13:27

EAST Search History

L33	0	re-organiz\$7 near (search near results)	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 13:28
L34	8	reorganiz\$7 near (search near results)	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 13:28
L35	5	34 and @ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 13:46
L36	4	35 and order\$4	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 13:44
L37	339	"latent semantic indexing"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 13:44
L38	10	reordering.near search near results	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 13:46
L39	2	38 and @ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 14:01
L40	6511	refresh\$3.ti.	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 14:00

EAST Search History

L41	7	refresh\$3.ti. and (search).ti.	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 14:00
L42	7	41 and @ad<"20031231"	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 14:01

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	8284	user near interface and ranking	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 15:49
L2	817	user near interface and re-ordering	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 15:49
L3	5	user near interface and re-ordering near search near results	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	OR	ON	2007/05/20 15:49

PORTAL
USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide

comparing sort order and search results

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used comparing sort order and search results Found 142,219 of 201,062

Sort results by relevance Save results to a Binder
Display results expanded form Search Tips Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: 1 2 3 4 5 6 7 8 9 10 next

Best 200 shown Relevance scale 

1 Inverted files for text search engines 

Justin Zobel, Alistair Moffat
July 2006 **ACM Computing Surveys (CSUR)**, Volume 38 Issue 2

Publisher: ACM Press
Full text available:  pdf(944.29 KB) Additional Information: full citation, abstract, references, index terms

The technology underlying text search engines has advanced dramatically in the past decade. The development of a family of new index representations has led to a wide range of innovations in index storage, index construction, and query evaluation. While some of these developments have been consolidated in textbooks, many specific techniques are not widely known or the textbook descriptions are out of date. In this tutorial, we introduce the key techniques in the area, describing both a core impl ...

Keywords: Inverted file indexing, Web search engine, document database, information retrieval, text retrieval

2 Comparative study of monolingual and multilingual search models for use with asian languages 

Jacques Savoy
June 2005 **ACM Transactions on Asian Language Information Processing (TALIP)**, Volume 4 Issue 2

Publisher: ACM Press
Full text available:  pdf(292.50 KB) Additional Information: full citation, abstract, references, index terms

Based on the NTCIR-4 test-collection, our first objective is to present an overview of the retrieval effectiveness of nine vector-space and two probabilistic models that perform monolingual searches in the Chinese, Japanese, Korean, and English languages. Our second goal is to analyze the relative merits of the various automated and freely available tools to translate the English-language topics into Chinese, Japanese, or Korean, and then submit the resultant query in order to retrieve pertinent ...

Keywords: Chinese language, Japanese language, Korean language, Multilingual information retrieval, cross-language information retrieval, natural language processing with Asian languages, results-merging, search engines with Asian languages

3 1 - Regular Articles: Fast string sorting using order-preserving compression 

PORTAL
USPTO

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search: The ACM Digital Library The Guide

search engine and sort order

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used [search engine](#) and [sort order](#) Found 77,303 of 201,062

Sort results by [relevance](#) Save results to a Binder [? Search Tips](#) Try an [Advanced Search](#)
Display results [expanded form](#) Open results in a new window Try this search in [The ACM Guide](#)

Results 1 - 20 of 200 Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#) Relevance scale 

1 [Inverted files for text search engines](#) 
 Justin Zobel, Alistair Moffat
 July 2006 **ACM Computing Surveys (CSUR)**, Volume 38 Issue 2
Publisher: ACM Press
 Full text available:  pdf(944.29 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)
 The technology underlying text search engines has advanced dramatically in the past decade. The development of a family of new index representations has led to a wide range of innovations in index storage, index construction, and query evaluation. While some of these developments have been consolidated in textbooks, many specific techniques are not widely known or the textbook descriptions are out of date. In this tutorial, we introduce the key techniques in the area, describing both a core impl ...

Keywords: Inverted file indexing, Web search engine, document database, information retrieval, text retrieval

2 [Computing customized page ranks](#) 
 Ah Chung Tsoi, Markus Hagenbuchner, Franco Scarselli
 November 2006 **ACM Transactions on Internet Technology (TOIT)**, Volume 6 Issue 4
Publisher: ACM Press
 Full text available:  pdf(1.63 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)
 In this article, we present a new approach to page ranking. The page rank of a collection of Web pages can be represented in a parameterized model, and the user requirements can be represented by a set of constraints. For a particular parameterization, namely, a linear combination of the page ranks produced by different forcing functions, and user requirements represented by a set of linear constraints, the problem can be solved using a quadratic programming method. The solution to this problem ...

Keywords: Interface personalization, PageRank, Web page scoring systems, search engines

3 [Indexing and querying: Sampling search-engine results](#) 
 Aris Anagnostopoulos, Andrei Z. Broder, David Carmel
 May 2005 **Proceedings of the 14th international conference on World Wide Web**
 WWW '05

PORTAL
USPTO

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search: The ACM Digital Library The Guide

search engine and re-sort

THE ACM DIGITAL LIBRARY

 [Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used [search engine](#) and [re sort](#)

Found 13,298 of 201,062

Sort results by

relevance

 [Save results to a Binder](#)

[Try an Advanced Search](#)

Display results

expanded form

 [Search Tips](#)

[Try this search in The ACM Guide](#)

Open results in a new window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale **1 Evaluation: Evaluation by comparing result sets in context**

 Paul Thomas, David Hawking

November 2006 **Proceedings of the 15th ACM international conference on Information and knowledge management CIKM '06**

Publisher: ACM Press

Full text available:  [pdf\(444.50 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Familiar evaluation methodologies for information retrieval (IR) are not well suited to the task of comparing systems in many real settings. These systems and evaluation methods must support contextual, interactive retrieval over changing, heterogeneous data collections, including private and confidential information. We have implemented a comparison tool which can be inserted into the natural IR process. It provides a familiar search interface, presents a small number of result sets in side-by-s ...

Keywords: embedded comparisons, evaluation

2 Web search 1: Using ODP metadata to personalize search

 Paul Alexandru Chirita, Wolfgang Nejdl, Raluca Paiu, Christian Kohlschütter

August 2005 **Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval SIGIR '05**

Publisher: ACM Press

Full text available:  [pdf\(310.29 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The Open Directory Project is clearly one of the largest collaborative efforts to manually annotate web pages. This effort involves over 65,000 editors and resulted in metadata specifying topic and importance for more than 4 million web pages. Still, given that this number is just about 0.05 percent of the Web pages indexed by Google, is this effort enough to make a difference? In this paper we discuss how these metadata can be exploited to achieve high quality personalized web search. First, we ...

Keywords: biased pageRank, metadata, open directory, personalized search

3 Posters: A light weight PDA-friendly collection fusion technique

 Jeffery Antoniuk, Mario A. Nascimento

July 2003 **Proceedings of the 26th annual international ACM SIGIR conference on Research and development in information retrieval SIGIR '03**


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
Search: The ACM Digital Library The Guide


THE ACM DIGITAL LIBRARY
[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used [blending search results](#)

Found 62,219 of 201,062

Sort results by

 [Save results to a Binder](#)

Display results

 [Search Tips](#)
 [Open results in a new window](#)
[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

Results 1 - 20 of 200

 Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

1 [Animation and motion: Quick transitions with cached multi-way blends](#)

Leslie Ikemoto, Okan Arik, David Forsyth

 April 2007 **Proceedings of the 2007 symposium on Interactive 3D graphics and games SI3D '07**
Publisher: ACM Press

 Full text available: [pdf\(225.49 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

We describe a discriminative method for distinguishing natural-looking from unnatural-looking motion. Our method is based on physical and data-driven features of motion to which humans seem sensitive. We demonstrate that our technique is significantly more accurate than current alternatives.

We use this technique as the testing part of a hypothesize-and-test motion synthesis procedure. The mechanism we build using this procedure can quickly provide an application with a transition of u ...

Keywords: motion blending, motion evaluation, motion synthesis

2 [Paper session IR-13 \(information retrieval\): context and personalization: Y!Q:](#)

[contextual search at the point of inspiration](#)

Reiner Kraft, Farzin Maghoul, Chi Chao Chang

 October 2005 **Proceedings of the 14th ACM international conference on Information and knowledge management CIKM '05**
Publisher: ACM Press

 Full text available: [pdf\(370.43 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Contextual search tries to better capture a user's information need by augmenting the user's query with contextual information extracted from the search context (for example, terms from the web page the user is currently reading or a file the user is currently editing). This paper presents Y!Q---a first of its kind large-scale contextual search system---and provides an overview of its system design and architecture. Y!Q solves two major problems. First, how to capture high quality search context. ...

Keywords: Web information retrieval, content analysis, context, contextual search, search personalization, semantic networks

PORTAL
USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide

search results and refresh or refreshing

THE ACM DIGITAL LIBRARY

 [Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used [search results and refresh or refreshing](#)

Found 61,364 of 201,062

Sort results by relevance [Save results to a Binder](#)

Try an [Advanced Search](#)

Display results expanded form [Search Tips](#)

Try this search in [The ACM Guide](#)

Open results in a new window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale **1 [User-focused search and crawling: User-centric Web crawling](#)**

 Sandeep Pandey, Christopher Olston

 May 2005 **Proceedings of the 14th international conference on World Wide Web**
WWW '05

Publisher: ACM Press

Full text available:  [pdf\(914.56 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Search engines are the primary gateways of information access on the Web today. Behind the scenes, search engines crawl the Web to populate a local indexed repository of Web pages, used to answer user search queries. In an aggregate sense, the Web is very dynamic, causing any repository of Web pages to become out of date over time, which in turn causes query answer quality to degrade. Given the considerable size, dynamicity, and degree of autonomy of the Web as a whole, it is not feasible for a ...

Keywords: Web crawling, Web page refreshing, user-centric

2 [Effective page refresh policies for Web crawlers](#)

 Junghoo Cho, Hector Garcia-Molina

 December 2003 **ACM Transactions on Database Systems (TODS)**, Volume 28 Issue 4

Publisher: ACM Press

Full text available:  [pdf\(345.52 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In this article, we study how we can maintain local copies of remote data sources "fresh," when the source data is updated autonomously and independently. In particular, we study the problem of *Web crawlers* that maintain local copies of remote Web pages for Web search engines. In this context, remote data sources (Websites) do not notify the copies (Web crawlers) of new changes, so we need to periodically *poll* the sources to maintain the copies up-to-date. Since polling the sources ...

Keywords: Web crawlers, page refresh, web search engines, world-wide web

3 [Paper session IR-13 \(information retrieval\): context and personalization: YIQ:](#)

 [contextual search at the point of inspiration](#)

 Reiner Kraft, Farzin Maghoul, Chi Chao Chang


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
Search: The ACM Digital Library The Guide



THE ACM DIGITAL LIBRARY

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used search results and refresh or refreshing and sort

Found 77,413 of 201,062

 Sort results by [Save results to a Binder](#)
[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

 Display results [Search Tips](#)
 [Open results in a new window](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

1 Searching the Web

Arvind Arasu, Jungwoo Cho, Hector Garcia-Molina, Andreas Paepcke, Sriram Raghavan
 August 2001 **ACM Transactions on Internet Technology (TOIT)**, Volume 1 Issue 1

Publisher: ACM Press
 Full text available: [pdf\(319.98 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

We offer an overview of current Web search engine design. After introducing a generic search engine architecture, we examine each engine component in turn. We cover crawling, local Web page storage, indexing, and the use of link analysis for boosting search performance. The most common design and implementation techniques for each of these components are presented. For this presentation we draw from the literature and from our own experimental search engine testbed. Emphasis is on introducing ...

Keywords: HITS, PageRank, authorities, crawling, indexing, information retrieval, link analysis, search engine

2 Skyline and similarity search: Refreshing the sky: the compressed skycube with efficient support for frequent updates

Tian Xia, Donghui Zhang
 June 2006 **Proceedings of the 2006 ACM SIGMOD international conference on Management of data SIGMOD '06**

Publisher: ACM Press
 Full text available: [pdf\(664.43 KB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

The skyline query is important in many applications such as multi-criteria decision making, data mining, and user-preference queries. Given a set of d-dimensional objects, the skyline query finds the objects that are not dominated by others. In practice, different users may be interested in different dimensions of the data, and issue queries on any subset of d dimensions. This paper focuses on supporting concurrent and unpredictable subspace skyline queries in frequently updated databases. Simpl ...

Keywords: compressed skycube, skyline, update scheme

3 Boosting the performance of Web search engines: Caching and prefetching query results by exploiting historical usage data

◆ Tiziano Fagni, Raffaele Perego, Fabrizio Silvestri, Salvatore Orlando
January 2006 **ACM Transactions on Information Systems (TOIS)**, Volume 24 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(668.69 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This article discusses efficiency and effectiveness issues in caching the results of queries submitted to a Web search engine (WSE). We propose SDC (Static Dynamic Cache), a new caching strategy aimed to efficiently exploit the temporal and spatial locality present in the stream of processed queries. SDC extracts from historical usage data the results of the most frequently submitted queries and stores them in a *static, read-only* portion of the cache. The remaining entries of the c ...

Keywords: Caching, Web search engines, multithreading

4 Information and control in gray-box systems



◆ Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau

October 2001 **ACM SIGOPS Operating Systems Review , Proceedings of the eighteenth ACM symposium on Operating systems principles SOSP '01**, Volume 35 Issue 5

Publisher: ACM Press

Full text available:  [pdf\(1.59 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In modern systems, developers are often unable to modify the underlying operating system. To build services in such an environment, we advocate the use of *gray-box* techniques. When treating the operating system as a gray-box, one recognizes that not changing the OS restricts, but does not completely obviate, both the *information* one can acquire about the internal state of the OS and the *control* one can impose on the OS. In this paper, we develop and investigate three gray-bo ...

5 Systems 3: searching and streaming: PRISM: indexing multi-dimensional data in P2P networks using reference vectors



◆ O. D. Sahin, A. Gulbeden, F. Emekci, D. Agrawal, A. El Abbadi

November 2005 **Proceedings of the 13th annual ACM international conference on Multimedia MULTIMEDIA '05**

Publisher: ACM Press

Full text available:  [pdf\(296.73 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Peer-to-peer (P2P) systems research has gained considerable attention recently with the increasing popularity of file sharing applications. Since these applications are used for sharing huge amounts of data, it is very important to efficiently locate the data of interest in such systems. However, these systems usually do not provide efficient search techniques. Existing systems offer only keyword search functionality through a centralized index or by query flooding. In this paper, we propose a s ...

Keywords: peer-to-peer systems, reference vectors, similarity search

6 Enhancing data warehouse performance through query caching



◆ Aditya N. Saharia, Yair M. Babad

May 2000 **ACM SIGMOD Database**, Volume 31 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(1.96 MB\)](#) Additional Information: [full citation](#), [index terms](#)

Keywords: adaptive query cache, database management, datawarehouse and repository, decision support, design, information systems applications, intelligent databases, management, performance, query minimization, query processing, query subsumption

7 Enhancing data warehouse performance through query caching



Aditya N. Saharia, Yair M. Babad
June 2000 **ACM SIGMIS Database**, Volume 31 Issue 3

Publisher: ACM Press

Full text available: [pdf\(1.96 MB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

The main function of a data warehouse is the separation of the decision layer from the operation layer so that users can invoke analysis, planning, and decision support applications without having to worry about constantly evolving operational databases. Such applications allow ad hoc queries for which no predefined reports exist. It is possible that an ad hoc query is submitted by different users or even by the same user at different times, requiring its repeated evaluations even though the con ...

Keywords: adaptive query cache, data warehouse and repository, database management, decision support, design, information systems applications, intelligent databases, management, performance, query minimization, query processing, query subsumption

8 Selected writings on computing: a personal perspective



Edsger W. Dijkstra
January 1982 Book

Publisher: Springer-Verlag New York, Inc.

Full text available: [pdf\(60.98 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#)

Since the summer of 1973, when I became a Burroughs Research Fellow, my life has been very different from what it had been before. The daily routine changed: instead of going to the University each day, where I used to spend most of my time in the company of others, I now went there only one day a week and was most of the time that is, when not travelling!-- alone in my study. In my solitude, mail and the written word in general became more and more important. The circumstance that my employe ...

9 Collected experience from implementing RSVP



Martin Karsten
August 2006 **IEEE/ACM Transactions on Networking (TON)**, Volume 14 Issue 4

Publisher: IEEE Press

Full text available: [pdf\(401.11 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Internet quality of service (QoS) is still a highly debated topic for more than fifteen years. Even with the large variety of QoS proposals and the impressive research advances, there is little deployment yet of network layer QoS technology. One specific problem domain is QoS signalling, which has recently attracted increasing attention to bring forward new standardization approaches. In this paper, an extensive study of RSVP is presented, covering protocol design, software design, and performan ...

Keywords: protocol implementation, quality of service, signalling

10

 Seeing, hearing, and touching: putting it all together 

Brian Fisher, Sidney Fels, Karon MacLean, Tamara Munzner, Ronald Rensink
August 2004 **ACM SIGGRAPH 2004 Course Notes SIGGRAPH '04**

Publisher: ACM Press

Full text available:  [pdf\(20.64 MB\)](#) Additional Information: [full citation](#)

11 Research sessions: Research 20: Reliability: Lazy database replication with snapshot isolation 

Khuzaima Daudjee, Kenneth Salem

September 2006 **Proceedings of the 32nd international conference on Very large data bases - Volume 32 VLDB'2006**

Publisher: VLDB Endowment

Full text available:  [pdf\(567.69 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Snapshot isolation is a popular transactional isolation level in database systems. Several replication techniques based on snapshot isolation have recently been proposed. These proposals, however, do not fully leverage the local concurrency controls that provide snapshot isolation. Furthermore, guaranteeing snapshot isolation in lazy replicated systems may result in transaction inversions, which happen when transactions see stale data. Strong snapshot isolation, which is provided in centralized ...

12 Research sessions: Research 28: Search applications: Automatic extraction of dynamic record sections from search engine result pages 

Hongkun Zhao, Weiyi Meng, Clement Yu

September 2006 **Proceedings of the 32nd international conference on Very large data bases - Volume 32 VLDB'2006**

Publisher: VLDB Endowment

Full text available:  [pdf\(897.74 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

A search engine returned result page may contain search results that are organized into multiple dynamically generated sections in response to a user query. Furthermore, such a result page often also contains information irrelevant to the query, such as information related to the hosting site of the search engine. In this paper, we present a method to automatically generate wrappers for extracting search result records from all dynamic sections on result pages returned by search engines. This me ...

13 Information Retrieval: Predictive caching and prefetching of query results in search engines 

 Ronny Lempel, Shlomo Moran

May 2003 **Proceedings of the 12th international conference on World Wide Web WWW '03**

Publisher: ACM Press

Full text available:  [pdf\(212.73 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We study the caching of query result pages in Web search engines. Popular search engines receive millions of queries per day, and efficient policies for caching query results may enable them to lower their response time and reduce their hardware requirements. We present PDC (probability driven cache), a novel scheme tailored for caching search results, that is based on a probabilistic model of search engine users. We then use a trace of over seven million queries submitted to the search engine A ...

Keywords: caching, query processing and optimization

14 Materialized view selection and maintenance using multi-query optimization

Hoshi Mistry, Prasan Roy, S. Sudarshan, Krithi Ramamritham

May 2001 **ACM SIGMOD Record , Proceedings of the 2001 ACM SIGMOD international conference on Management of data SIGMOD '01**, Volume 30 Issue 2**Publisher:** ACM PressFull text available: [pdf\(199.46 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Materialized views have been found to be very effective at speeding up queries, and are increasingly being supported by commercial databases and data warehouse systems. However, whereas the amount of data entering a warehouse and the number of materialized views are rapidly increasing, the time window available for maintaining materialized views is shrinking. These trends necessitate efficient techniques for the maintenance of materialized views.

In this paper, we show how to find an ...

15 Schemes of storing XML query cache

Hyunchul Kang, Seungchul Han, Younghyun Kim

January 2005 **Proceedings of the 16th Australasian database conference - Volume 39 ADC '05****Publisher:** Australian Computer Society, Inc.Full text available: [pdf\(542.38 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

XML query caching for XML database-backed Web applications began to be investigated recently. However, the issue of how the cached query results are stored has *not* been addressed despite its practical significance. In this paper, we deal with the schemes of *storing* XML query cache. A fundamental problem encountered in designing an efficient storage structure for XML query cache is that there exist performance tradeoffs between the two major types of operations on a cached query res ...

Keywords: DOM, PDOM, XML query cache, XML update, incremental cache refresh

16 GPGPU: general purpose computation on graphics hardware

David Luebke, Mark Harris, Jens Krüger, Tim Purcell, Naga Govindaraju, Ian Buck, Cliff Woolley, Aaron Lefohn

August 2004 **ACM SIGGRAPH 2004 Course Notes SIGGRAPH '04****Publisher:** ACM PressFull text available: [pdf\(63.03 MB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#)

The graphics processor (GPU) on today's commodity video cards has evolved into an extremely powerful and flexible processor. The latest graphics architectures provide tremendous memory bandwidth and computational horsepower, with fully programmable vertex and pixel processing units that support vector operations up to full IEEE floating point precision. High level languages have emerged for graphics hardware, making this computational power accessible. Architecturally, GPUs are highly parallel s ...

17 Status report of the graphic standards planning committee

Computer Graphics staff

August 1979 **ACM SIGGRAPH Computer Graphics**, Volume 13 Issue 3**Publisher:** ACM PressFull text available: [pdf\(15.01 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#)

◆ Maintenance of K-nn and spatial join queries on continuously moving points 

Glenn S. Iwerks, Hanan Samet, Kenneth P. Smith

June 2006 **ACM Transactions on Database Systems (TODS)**, Volume 31 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(2.49 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Cars, aircraft, mobile cell phones, ships, tanks, and mobile robots all have the common property that they are moving objects. A kinematic representation can be used to describe the location of these objects as a function of time. For example, a moving point can be represented by the function $p(t) = l_0 + (t - t_0) \cdot v$, where l_0 is the start location, t_0 is the st ...

Keywords: *k*-nearest neighbor, Moving object databases, continuously moving objects, materialized view maintenance, spatial join, temporal databases

19 Classics in software engineering 

January 1979 Divisible Book

Publisher: Yourdon Press

Full text available:  [pdf\(22.45 MB\)](#) Additional Information: [full citation](#), [cited by](#), [index terms](#)

20 Task & attention: InkSeine: In Situ search for active note taking 

◆ Ken Hinckley, Shengdong Zhao, Raman Sarin, Patrick Baudisch, Edward Cutrell, Michael Shilman, Desney Tan

April 2007 **Proceedings of the SIGCHI conference on Human factors in computing systems CHI '07**

Publisher: ACM Press

Full text available:  [pdf\(1.27 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Using a notebook to sketch designs, reflect on a topic, or capture and extend creative ideas are examples of active note taking tasks. Optimal experience for such tasks demands concentration without interruption. Yet active note taking may also require reference documents or emails from team members. InkSeine is a Tablet PC application that supports active note taking by coupling a pen-and-ink interface with an in situ search facility that flows directly from a user's ink notes (Fig. 1). InkS ...

Keywords: gestures, handwriting, ink, input, pen, search, tablet

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

[Search Session History](#)[BROWSE](#)[SEARCH](#)[IEEE Xplore GUIDE](#)

Sun, 20 May 2007, 3:38:05 PM EST

Edit an existing query or
compose a new query in the
Search Query Display.

Select a search number (#)
to:

- Add a query to the Search Query Display
- Combine search queries using AND, OR, or NOT
- Delete a search
- Run a search

Recent Search Queries#1 ((comparing sort order)<in>metadata)#2 (comparing and sort and order<IN>metadata)#3 (comparing and sort and order<in>metadata) and re-order#4 (comparing and search near order<IN>metadata)#5 (comparing and rank and order<IN>metadata)#6 (search and engine and search and results and ordering<IN>metadata)#7 (search and engine and search and results and ordering<in>metadata) and altavista#8 (search and engine and search and results and ordering<in>metadata) and altavista[Clear Session History](#)[Help](#) [Contact Us](#) [Privacy & :](#)

© Copyright 2006 IEEE -

Indexed by
 Inspec®